Finding Naked People

Margaret M. Fleck and David A Forsyth

The recent explosion in internet usage and multi-media computing has created a substantial demand for algorithms that perform **content-based retrieval**, i.e. that select images from a large database based on what they depict. Identifying images depicting **naked** or scantily-dressed **people** is a natural content-based retrieval problem. These images frequently lack textual labels adequate to identify their content but can be effectively detected using simple visual cues (color, texture, simple shape features), of the type that the human visual system is known to use for fast (preattentive) triage.

From a computer vision point of view, this problem is difficult because we have little control over pose or imaging conditions, and our objects are flexible. Images of **naked people** found on the internet may have any of a wide range of backgrounds. They often depict multiple figures or partial figures. Limbs may be arranged in many different positions. The images are taken from a wide variety of camera angles and under a wide range of lighting conditions.

We have implemented an algorithm that finds images of **naked people** without human assistance. This algorithm analyzes images in two steps:

- **Skin Filter**: The algorithm first locates images containing large areas whose color and texture is appropriate for skin,
- **Grouper**: Within these areas, the algorithm finds elongated regions and groups them into possible human limbs and connected groups of limbs, using specialised groupers which incorporate substantial amounts of information about object structure.

Images containing sufficiently large skin-colored groups of possible limbs are reported as potentially containing **naked people**.

This algorithm was tested on a database of 4854 images: 565 images of **naked people** and 4289 control images from a variety of sources. The skin filter identified 448 test images and 485 control images as containing substantial areas of skin. Of these, the grouper identified 241 test images and 182 control images as containing **people**-like shapes.

Thus, this algorithm correctly found 43% of the test images in the database. Furthermore, the test images represented 57% of its output. This would be extremely good performance for any computer vision or

database algorithm. It is unprecedented for a content-based query on uncontrolled image data.

For more information on our algorithm, see:

- Margaret Fleck, David Forsyth, and Chris Bregler (1996) ``**Finding Naked People**," 1996 **European Conference on Computer Vision** , Volume II, pp. 592-602. (postscript)
- David A. Forsyth and Margaret M. Fleck (1996) ``Identifying nude pictures," IEEE Workshop on the Applications of Computer Vision 1996, pp. 103--108. (postscript)
- The on-line details of our skin filter

See also more recent work on recognizing horses using automatically-learned body plans.

This research was carried out at the University of Iowa and at the University of California, Berkeley.

This page is maintained by Margaret Fleck.